
# Netfilter development updates in 2015

Pablo Neira Ayuso
<pablo@netfilter.org>

Linux kernel meetup – Ankert
June 2015
Budapest, Hungary

# nftables?

- What is?
  - Replacement for {ip,ip6,arp,eb}tables.
- Why?
  - Avoid code/tool duplication and inconsistencies.
  - Faster packet classification through enhanced generic set infrastructure.
  - Simplified dual stack IPv4/IPv6 administration.
  - Better userspace API.
  - Nice syntax.

# Tables, chains and rules

- nft add table ip foo
- nft add chain ip foo bar { \
  type filter hook input priority 0; policy drop; \
  }
- nft add rule ip foo bar \
  ct state established,related accept
  nft add rule ip foo bar \
  ct state new tcp dport 22 accept

# Expressions

- nft add rule ip foo bar tcp dport != 80
- nft add rule ip foo bar tcp dport 1-1024
- nft add rule ip foo bar meta skuid 1000-1100
- nft add rule ip foo bar ip daddr 192.168.10.0/24
- nft add rule ip foo bar meta mark 0xffffffff/24
- nft add rule ip foo bar ct  state new,established
- nft add rule ip foo bar ct mark and 0x0000ffff == 0x0000123
- nft add rule ip foo bar ct mark set 10
- nft add rule ip foo bar ct mark set meta mark

# Sets and maps

- nft add rule ip foo bar tcp dport { 22, 80, 443 } counter
- nft add set ip foo whitelist { type ipv4_addr \; }
  nft add rule ip foo bar ip daddr @whitelist counter accept
  nft add element ip foo whitelist { \
  192.168.0.1, \
  192.168.0.10 \
  }
- nft add table ip nat
  nft add chain ip nat post { \
  type nat hook postrouting priority 0\; }
  nft add rule ip nat post snat ip saddr map { \
  1.1.1.0/24 : 192.168.3.11 , \
  2.2.2.0/24 : 192.168.3.12 \
  }

# Dictionaries

- nft add chain ip foo tcp-chain
  nft add chain ip foo udp-chain
  nft add chain ip foo icmp-chain

- nft add rule ip foo bar ip protocol vmap { \
  tcp : jump tcp-chain, \
  udp : jump udp-chain, \
  icmp : jump icmp-chain
  }

# What's new in the next release?

- nft add rule ip foo bar ip saddr . tcp dport { \
  192.168.1.123 . 22 : accept, \
  192.168.1.123 . 80 : accept, \
  }
- nft add set ip foo whitelist { \
  type ipv4_addr; \
  timeout 1h; \
  }
- nft add element ip foo whitelist { \
  192.168.10.123 comment “temporary access” \
  }

# Learn more and help us

- Grab the code
  - Kernel: http://www.kernel.org
  - Library: git://git.netfilter.org/libnftnl
  - User-space: git://git.netfilter.org/nftables
- Documentation
  - http://wiki.nftables.org
  - man nft
- Report bugs:
  - https://bugzilla.netfilter.org


# Netfilter development updates in 2015

Pablo Neira Ayuso
<pablo@netfilter.org>

Linux kernel meetup – Ankert
June 2015
Budapest, Hungary